arXiv:1605.08278v4 [cs.SE] 20 Mar 2017

Should We Learn Probabilistic Models for Model Checking? A New Approach and an Empirical Study*

Jingyi Wang[1], Jun Sun[1], Qixia Yuan[2] and Jun Pang[2]

[1] Singapore University of Technology and Design
{jingyi_wang@mymail.,sunjun@}sutd.edu.sg
[2] University of Luxembourg
{qixia.yuan,jun.pang}@uni.lu

Abstract. Many automated system analysis techniques (e.g., model checking, model-based testing) rely on first obtaining a model of the system under analysis. System modeling is often done manually, which is often considered as a hindrance to adopt model-based system analysis and development techniques. To overcome this problem, researchers have proposed to automatically "learn" models based on sample system executions and shown that the learned models can be useful sometimes. There are however many questions to be answered. For instance, how much shall we generalize from the observed samples and how fast would learning converge? Or, would the analysis result based on the learned model be more accurate than the estimation we could have obtained by sampling many system executions within the same amount of time? In this work, we investigate existing algorithms for learning probabilistic models for model checking, propose an evolution-based approach for better controlling the degree of generalization and conduct an empirical study in order to answer the questions. One of our findings is that the effectiveness of learning may sometimes be limited.

Keywords: probabilistic model checking, model learning, genetic algorithm

1 Introduction

Many system analysis techniques rely on first obtaining a system model. The model should be accurate and often is required to be at a proper level of abstraction. For instance, model checking [10,3] works effectively if the user-provided model captures all the relevant behavior of the system and abstracts away the irrelevant details. With such a model as well as a given property, a model checker would automatically verify the property or falsify it with a counterexample. Alternatively, in the setting of probabilistic model checking (PMC, see Sect. 2) [3,5], the model checker would calculate the probability of satisfying the property.

Model checking is perhaps not as popular as it ought to be due to the fact that a good model is required beforehand. For instance, a model which is too general would introduce spurious counterexamples, whereas the checking result based on a model which under-approximates the relevant system behavior is untrustworthy. In the setting

* This work was supported by NRF Award No. NRF2014NCR-NCR001-40.

of PMC, users are required to provide a probabilistic model (e.g., a Markov chain [3]) with accurate probabilistic distributions, which is often challenging.

In practice, system modeling is often done manually, which is both time-consuming and error-prone. Worse, it could be infeasible if the system is a black box or it is so complicated that no accurate model is known (e.g., the chemical reaction in a water treatment system [35]). This is often considered by industry as one hindrance to adopt otherwise powerful techniques like model checking. Alternative approaches which would rely less on manual modeling have been explored in different settings. One example is statistical model checking (SMC, see Sect. 2) [41,33]. The main idea is to provide a statistical measure on the likelihood of satisfying a property, by observing sample system executions and applying standard techniques like hypothesis testing [4,13,41]. SMC is considered useful partly because it can be applied to black-box or complex systems when system models are not available.

Another approach for avoiding manual modeling is to automatically learn models. A variety of learning algorithms have been proposed to learn a variety of models, e.g., [32,31,7,17]. It has been showed that the learned models can be useful for subsequent system analysis in certain settings, especially so when having a model is a must. Recently, the idea of model learning has been extended to system analysis through model checking. In [23,9,24], it is proposed to learn a probabilistic model first and then apply techniques like PMC to calculate the probability of satisfying a property based on the learned model. On one hand, learning is beneficial, and it solves some known drawbacks of SMC or even simulation-based system analysis methods in general. For instance, since SMC relies on sampling *finite* system executions, it is challenging to verify un-bounded properties [39,29], whereas we can verify un-bounded properties based on the learned model through PMC. Furthermore, the learned model can be used to facilitate other system analysis tasks like model-based testing and software simulation for complicated systems. On the other hand, learning essentially is a way of generalizing the sample executions and there are often many variables. It is thus worth investigating how the sample executions are generalized and whether indeed such learning-based approaches are justified.

In particular, we would like to investigate the following research questions. Firstly, how can we control the degree of generalization for the best learning outcome, since it is known that both over-fitting or under-fitting would cause problems in subsequent analysis? Secondly, often it is promised that the learned model would converge to an accurate model of the original system, if the number of sample executions is sufficiently large. In practice, there could be only a limited number of sample executions and thus it is valid to question how fast the learning algorithms converge. Furthermore, do learning-based approaches offer better analysis results if alternative approaches which do not require a learned model, like SMC, are available?

In order to answer the above questions, we mainly make the following contributions. Firstly, we propose a new approach (Sect. 4) to better control the degree of generalization than existing approaches (Sect. 3) in model learning. The approach is inspired by our observations on the limitations of existing learning approaches. Experiment results show that our approach converges faster than existing approaches while providing better or similar analysis results. Secondly, we develop a software toolkit ZIQIAN, realizing

previously proposed learning approaches for PMC as well as our approach so as to systematically study and compare them in a fair way. Lastly, we conduct an empirical study on comparing different model learning approaches against a suite of benchmark systems, two real world systems, as well as randomly generated models (Sect. 5). One of our findings suggests that learning models for model checking might not be as effective compared to SMC given the same time limit. However, the learned models may be useful when manual modeling is impossible. From a broader point of view, our work is a first step towards investigating the recent trend on adopting machine learning techniques to solve software engineering problems. We remark there are extensive existing research on learning non-probabilistic models (e.g., [1]), which is often designed for different usage and is thus beyond the scope of this work. We review related work and conclude this paper in Sect. 6.

2 Preliminary

In this work, the model that we focus on is discrete-time Markov chains (DTMC) [3]. The reason is that most existing learning algorithms generate DTMC and it is still ongoing research on how to learn other probabilistic models like Markov Decision Processes [6,24,32,23,9]. Furthermore, the learned DTMC is aimed for probabilistic analysis by methods like PMC, among others. In the following, we briefly introduce DTMC, PMC as well as SMC so that we can better understand the context.

Markov Chain A DTMC $\mathcal{D}$ is a triple tuple (S, ι_{init}, Tr), where S is a countable, nonempty set of states; $\iota_{init} : S \to [0, 1]$ is the initial distribution s.t. $\sum_{s \in S} \iota_{init}(s) = 1$; and $Tr : S \times S \to [0, 1]$ is the transition probability assigned to every pair of states which satisfies the following condition: $\sum_{s' \in S} Tr(s, s') = 1$. $\mathcal{D}$ is *finite* if S is finite. For instance, an example DTMC modelling the *egl* protocol [21] is shown in Figure 1.

A DTMC induces an underlying digraph where states are vertices and there is an edge from s to s' if and only if $Tr(s, s') > 0$. *Paths* of DTMCs are maximal paths in the underlying digraph, defined as infinite state sequences $\pi = s_0 s_1 s_2 \cdots \in S^\omega$ such that $Tr(s_i, s_{i+1}) > 0$ for all $i \geq 0$. We write $Path^{\mathcal{D}}(s)$ to denote the set of all infinite paths of $\mathcal{D}$ starting from state s.

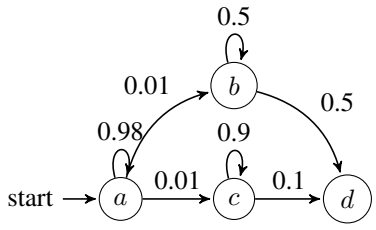

Fig. 1: DTMC of *egl* protocol.

Probabilistic Model Checking PMC [5,3] is a formal analysis technique for stochastic systems including DTMCs. Given a DTMC $\mathcal{D} = (S, \iota_{init}, Tr)$ and a set of propositions Σ, we can define a function $L : S \to \Sigma$ which assigns valuation of the propositions in Σ to each state in S. Once each state is labeled, given a path in $Path^{\mathcal{D}}(s)$, we can obtain a corresponding sequence of propositions labeling the states.

Let $\Sigma^\star$ and Σ^ω be the set of all finite and infinite strings over Σ respectively. A property of the DTMC can be specified in temporal logic. Without loss of generality,

we focus on Linear Time Temporal logic (LTL) and probabilistic LTL in this work. An LTL formula φ over Σ is defined by the syntax:

$$\varphi ::= true \mid \sigma \mid \varphi_1 \wedge \varphi_2 \mid \neg\varphi \mid \mathbf{X}\varphi \mid \varphi_1 \mathbf{U} \varphi_2$$

where $\sigma \in \Sigma$ is a proposition; **X** is intuitively read as 'next' and **U** is read as 'until'. We remark commonly used temporal operators like **F** (which reads 'eventually') and **G** (which reads 'always') can be defined using the above syntax, e.g., $\mathbf{F}\varphi$ is defined as $true\mathbf{U}\varphi$. Given a string π in $\Sigma^\star$ or Σ^ω, we define whether π satisfies a given LTL formula φ in the standard way [3].

Given a path π of a DTMC, we write $\pi \models \varphi$ to denote that the sequence of propositions obtained from π satisfies φ and $\pi \not\models \varphi$ otherwise. Furthermore, a probabilistic LTL formula ϕ of the form $Pr_{\bowtie r}(\varphi)$ can be used to quantify the probability of a system satisfying the LTL formula φ, where $\bowtie \in \{\geq, \leq, =\}$ and $r \in [0, 1]$ is a probability threshold. A DTMC $\mathcal{D}$ satisfies $Pr_{\bowtie r}(\varphi)$ if and only if the accumulated probability of all paths obtained from the initial state of $\mathcal{D}$ which satisfy φ satisfies the condition $\bowtie r$. Given a DTMC $\mathcal{D}$ and a probabilistic LTL property $Pr_{\bowtie r}(\varphi)$, the PMC problem can be solved using methods like the automata-theoretic approach [3]. We skip the details of the approach and instead remark that the complexity of PMC is doubly exponential in the size of φ and polynomial in the size of $\mathcal{D}$.

Statistical Model Checking SMC is a Monte Carlo method to solve the probabilistic verification problem based on system simulations. Its biggest advantage is perhaps that it does not require the availability of system models [11]. SMC works by sampling system behaviors randomly (according to certain underlying probabilistic distribution) and observing how often a given property φ is satisfied. The idea is to provide a statistical measure on the likelihood of satisfying φ based on the observations, by applying techniques like hypothesis testing [4,13,41]. We refer readers to [41,3] for details.

3 Probabilistic Model Learning

Learning models from sample system executions for the purpose of PMC has been explored extensively in recent years [32,31,7,17,23,9,24]. In this section, we briefly present existing model learning algorithms for two different settings.

3.1 Learn from Multiple Executions

In the setting that the system can be reset and restarted multiple times, a set of independent executions of the system can be collected as input for learning. Learning algorithms in this category make the following assumptions [23]. First, the underlying system can be modeled as a DTMC. Second, the sampled system executions are mutually independent. Third, the length of each simulation is independent.

Let Σ denote the alphabet of the system observations such that each letter $e \in \Sigma$ is an observation of the system state. A system execution is then a finite string over Σ. The input in this setting is a finite set of strings $\Pi \subseteq \Sigma^\star$. For any string $\pi \in \Sigma^\star$, let $prefix(\pi)$ be the set of all prefixes of π including the empty string $\langle\rangle$. Let $prefix(\Pi)$ be the set of all prefixes of any string $\pi \in \Pi$. The set of strings Π can be naturally organized into a tree $tree(\Pi) = (N, root, E)$ where each node in N is a member of

$prefix(\Pi)$; the root is the empty string $\langle\rangle$; and $E \subseteq N \times N$ is a set of edges such that (π, π') is in E if and only if there exists $e \in \Sigma$ such that $\pi \cdot \langle e \rangle = \pi'$ where $\cdot$ is the sequence concatenation operator.

The idea of the learning algorithms is to generalize $tree(\Pi)$ by merging the nodes according to certain criteria in certain fixed order. Intuitively, two nodes should be merged if they are likely to represent the same state in the underlying DTMC. Since we do not know the underlying DTMC, whether two states should be merged is decided through a procedure called *compatibility test*. We remark the compatibility test effectively controls the degree of generalization. Different types of compatibility test have been studied [7,30,20]. We present in detail the compatibility test adopted in the AALERGIA algorithm [23] as a representative. First, each node π in $tree(\Pi)$ is labeled with the number of strings str in Π such that π is a prefix of str. Let $L(\pi)$ denote its label. Two nodes π_1 and π_2 in $tree(\Pi)$ are considered compatible if and only if they satisfy two conditions. The first condition is $last(\pi_1) = last(\pi_2)$ where $last(\pi)$ is the last letter in a string π, i.e., if the two nodes are to be merged, they must agree on the last observation (of the system state). The second condition is that the future behaviors from π_1 and π_2 must be sufficiently similar (i.e., within Angluin's bound [2]). Formally, given a node π in $tree(\Pi)$, we can obtain a probabilistic distribution of the next observation by *normalizing* the labels of the node and its children. In particular, for any event $e \in \Sigma$, the probability of going from node π to $\pi \cdot \langle e \rangle$ is defined as: $Pr(\pi, \langle e \rangle) = \frac{L(\pi \cdot \langle e \rangle)}{L(\pi)}$. We remark the probability of going from node π to itself is $Pr(\pi, \langle\rangle) = 1 - \sum_{e \in \Sigma} Pr(\pi, \langle e \rangle)$, i.e., the probability of not making any more observation. The multi-step probability from node π to $\pi \cdot \pi'$ where $\pi' = \langle e_1, e_2, \cdots, e_k \rangle$, written as $Pr(\pi, \pi')$, is the product of the one-step probabilities:

$$Pr(\pi, \pi') = Pr(\pi, \langle e_1 \rangle) \times Pr(\pi \cdot \langle e_1 \rangle, \langle e_2 \rangle) \times \cdots \times Pr(\pi \cdot \langle e_1, e_2, \cdots, e_{k-1} \rangle, \langle e_k \rangle) \tag{1}$$

Two nodes π_1 and π_2 are compatible if the following is satisfied:

$$|Pr(\pi_1, \pi) - Pr(\pi_2, \pi)| < \sqrt{6\epsilon \log(L(\pi_1))/L(\pi_1)} + \sqrt{6\epsilon \log(L(\pi_2))/L(\pi_2)} \tag{2}$$

for all $\pi \in \Sigma^\star$. We highlight that ϵ used in the above condition is a parameter which effectively controls the degree of state merging. Intuitively, a larger ϵ leads to more state merging, thus fewer states in the learned model.

If π_1 and π_2 are compatible, the two nodes are merged, i.e., the tree is transformed such that the incoming edge of π_2 is directed to π_1. Next, for any $\pi \in \Sigma^*$, $L(\pi_1 \cdot \pi)$ is incremented by $L(\pi_2 \cdot \pi)$. The algorithm works by iteratively identifying nodes which are compatible and merging them until there are no more compatible nodes. After merging all compatible nodes, the last phase of the learning algorithms in this category is to normalize the tree so that it becomes a DTMC.

3.2 Learn from a Single Execution

In the setting that the system cannot be easily restarted, e.g., real-world cyber-physical systems. We are limited to observe the system for a long time and collect a single, long execution as input. Thus, the goal is to learn a model describing the long-run, stationary

behavior of a system, in which system behaviors are decided by their finite variable length memory of the past behaviors.

In the following, we fix α to be the single system execution. Given a string $\pi = \langle e_0, e_1, \cdots, e_k \rangle$, we write $\mathit{suffix}(\pi)$ to be the set of all suffixes of π, i.e., $\mathit{suffix}(\pi) = \{\langle e_i, \cdots, e_k \rangle | 0 \leq i \leq k\} \cup \{\langle\rangle\}$. Learning algorithms in this category [9,31] similarly construct a tree $tree(\alpha) = (N, root, E)$ where N is the set of suffixes of α; $root = \langle\rangle$; and there is an edge $(\pi_1, \pi_2) \in E$ if and only if $\pi_2 = \langle e \rangle \cdot \pi_1$. For any string π, let $\#(\pi, \alpha)$ be the number of times π appears as a substring in α. A node π in $tree(\alpha)$ is associated with a function Pr_π such that $Pr_\pi(e) = \frac{\#(\pi \cdot \langle e \rangle, \alpha)}{\#(\pi, \alpha)}$ for every $e \in \Sigma$, which is the likelihood of observing e next given the previous observations π. Effectively, function Pr_π defines a probabilistic distribution of the next observation.

Based on different suffixes of the execution, different probabilistic distributions of the next observation will be formed. For instance, the probabilistic distribution from the node $\langle e \rangle$ where e is the last observation would predict the distribution only based on the last observation, whereas the node corresponding to the sequence of all previous observations would have a prediction based the entire history. The central question is how far we should look into the past in order to predict the future. As we observe more history, we will make a better prediction of the next observation. Nonetheless, constructing the tree completely (no generalization) is infeasible and the goal of the learning algorithms is thus to grow a part of the tree which would give a "good enough" prediction by looking at a small amount of history. The questions are then: what is considered "good enough" and how much history is necessary. The answers control the degree of generalization in the learned model.

In the following, we present the approach in [9] as a representative of algorithms proposed in the setting. Let $\mathit{fre}(\pi, \alpha) = \frac{\#(\pi, \alpha)}{|\alpha| - |\pi| - 1}$ where $|\pi|$ is the length of π be the relative frequency of having substring π in α. Algorithm 1 shows the algorithm for identifying the right tree by growing it on-the-fly. Initially, at line 1, the tree T contains only the root $\langle\rangle$. Given a threshold ϵ, we identify the set $S = \{\pi | \mathit{fre}(\pi, \alpha) > \epsilon\}$ at line 2, which are substrings appearing often enough in α and are candidate nodes to grow in the tree. The loop from line 3 to 7 keeps growing T. In particular, given a candidate node π, we find the longest suffix π' in T at line 4 and if we find that adding π would improve the prediction of the next observations by at least ϵ, π is added, along with all of its suffixes if they are currently missing from the tree (so that we maintain all suffixes of all nodes in the tree all the time). Whether we add node π into tree T or not, we update the set of candidate S to include longer substrings of α at line 6. When Algorithm 1 terminates, the tree contains all nodes which would make a *good enough* prediction. Afterwards, the tree is transformed into a DTMC where the leafs of $tree(\alpha)$ are turned into states in the DTMC (refer to [31] for details).

4 Learning through Evolution

Model learning essentially works by generalizing the sample executions. The central question is thus how to control the degree of generalization. To find the best degree of generalization, both [23] and [9] proposed to select the 'optimal' ϵ value using the golden section search of the highest Bayesian Information Criterion (BIC) score. For

Algorithm 1 $Learn\ PST$

1: Initialize T to be a single root node representing $\langle\rangle$;
2: Let $S = \{\sigma | fre(\sigma, \alpha) > \epsilon\}$ be the candidate suffix set;
3: **while** S is not empty **do**
4: Take any π from S; Let π' be the longest suffix of π in T;
5: (B) If $fre(\pi, \alpha) \cdot \sum_{\sigma \in \Sigma} Pr(\pi, \sigma) \cdot \log \frac{Pr(\pi, \sigma)}{Pr(\pi', \sigma)} \geq \epsilon$
add π and all its suffixes which are not in T to T;
6: (C) If $fre(\pi, \alpha) > \epsilon$, add $\langle e \rangle \cdot \pi$ to $\mathcal{S}$ for every $e \in \Sigma$ if $fre(\langle e \rangle \cdot \pi, \alpha) > 0$;
7: **end while**

instance, in [23], the BIC score of a learned model M, given the sample executions Π, is computed as follows: $log(Pr_M(\Pi)) - \mu \times |M| \times log(|\Pi|)$ where $|M|$ is the number of states in M; Π is the total number of observations and μ is a constant (set to be 0.5 in [23]) which controls the relative importance of the size of the learned model. This kind of approach to optimize BIC is based on the assumption that the BIC score is a concave function of the parameter ϵ. Our empirical study (refer to details in section 5), however, shows that this assumption is flawed and the BIC score can fluctuate with ϵ.

In the following, we propose an alternative method for learning models based on genetic algorithms (GA) [18]. The method is designed to select the best degree of generalization without the assumption of BIC's concaveness. The idea is that instead of using a predefined ϵ value to control the degree of generalization, we systematically generate candidate models and select the ones using the principle of natural selection so that the "fittest" model is selected eventually. In the following, we first briefly introduce the relevant background on GA and then present our approach in detail.

4.1 Genetic Algorithms

GA [18] are a set of optimization algorithms inspired by the "survival of the fittest" principle of Darwinian theory of natural selection. Given a specific problem whose solution can be encoded as a chromosome, a genetic algorithm typically works in the following steps [12]. First, an initial population (i.e., candidate solutions) is created either randomly or hand-picked based on certain criteria. Second, each candidate is evaluated using a pre-defined fitness function to see how good it is. Third, those candidates with higher fitness scores are selected as the parents of the next generation. Fourth, a new generation is generated by genetic operators, which either randomly alter (a.k.a. mutation) or combine fragments of their parent candidates (a.k.a. cross-over). Lastly, step 2-4 are repeated until a satisfactory solution is found or some other termination condition (e.g., timeout) is satisfied. GA are especially useful in providing approximate 'optimal' solutions when other optimization techniques do not apply or are too expensive, or the problem space is too large or complex.

GA are suitable for solving our problem of learning DTMC because we view the problem as finding an optimal DTMC model which not only maximizes the likelihood of the observed system executions but also satisfies additional constrains like having a small number of states. To apply GA to solve our problem, we need to develop a way of encoding candidate models in the form of chromosomes, define operators such

as mutation and crossover to generate new candidate models, and define the fitness function to selection better models. In the following, we present the details of the steps in our approach.

4.2 Learn from Multiple Executions

We first consider the setting where multiple system executions are available. Recall that in this setting, we are given a set of strings Π, from which we can build a tree representation $tree(\Pi)$. Furthermore, a model is learned through merging the nodes in $tree(\Pi)$. The space of different ways of merging the nodes thus corresponds to the potential models to learn. Our goal is to apply GA to search for the best model in this space. In the following, we first show how to encode different ways of merging the nodes as chromosomes.

Let the size of $tree(\Pi)$ (i.e., the number of nodes) be X and let Z be the number of states in the learned model. A way of merging the nodes is a function which maps each node in $tree(\Pi)$ to a state in the learned model. That is, it can be encoded as a chromosome in the form of a sequence of integers $\langle I_1, I_2, \cdots, I_X \rangle$ where $1 \leq I_i \leq Z$ for all i such that $1 \leq i \leq X$. Intuitively, the number I_i means that node i in $tree(\Pi)$ is mapped into state I_i in the learned model. Besides, the encoding is done such that infeasible models are always avoided. Recall that two nodes π_1 and π_2 can be merged only if $last(\pi_1) = last(\pi_2)$, which means that two nodes with different last observation should not be mapped into the same state in the learned model. Thus, we first partition the nodes into $|\Sigma|$ groups so that all nodes sharing the same last observation are mapped to the same group of integers. A chromosome is then generated such that only nodes in the same group can possibly be mapped into the same state. The initial population is generated by randomly generating a set of chromosomes this way. *We remark that in this way all generated chromosomes represent a valid DTMC model.*

Formally, the chromosome $\langle I_1, I_2, \cdots, I_X \rangle$ represents a DTMC $M = (S, \iota_{init}, Tr)$ where S is a set of Z states. Each state s in S corresponds to a set of nodes in $tree(\Pi)$. Let $nodes(s)$ denote that set. Tr is defined such that for all states s and s' in M,

$$Tr(s, s') = \frac{\sum_{x \in nodes(s)} \sum_{e \in \Sigma | \langle s,e \rangle \in nodes(s')} L(x \cdot \langle e \rangle)}{\sum_{x \in nodes(s)} L(x)} \tag{3}$$

The initial distributions ι_{init} is defined such that for any state $s \in S$, $\iota_{init}(s) = \sum_{x \in nodes(s)} L(x)/L(\langle\rangle)$.

Next, we define the fitness function. Intuitively, a chromosome is good if the corresponding DTMC model M maximizes the probability of the observed sample executions and the number of states in M is small. We thus define the fitness function of a chromosome as: $log(Pr_M(\Pi)) - \mu \times |M| \times log|\Pi|$ where $|M|$ is the number of states in M and $|\Pi|$ is the total number of letters in the observations and μ is a constant which represents how much we favor a smaller model size. The fitness function, in particular, the value of μ, controls the degree of generalization. If μ is 0, $tree(\Pi)$ would be the resultant model; whereas if μ is infinity, a model with one state would be generated. We remark that this fitness function is the same as the formula for computing the BIC score in [23]. Compared to existing learning algorithms, controlling the degree of generalization in our approach is more intuitive (i.e., different value of μ has a direct effect on

Algorithm 2 Model learning by GA from multiple executions

input: $tree(\Pi)$ and the alphabet Σ
output: A chromosome encoding a DTMC $\mathcal{D}$

Let Z be $|\Sigma|$; Let $Best$ be $null$;
repeat
  Let $population$ be an initial population with Z states;
  Let $generation$ be 1;
  repeat
    Let $newBest$ be the fittest in $population$;
    if $newBest$ is fitter than $Best$ **then**
      Set $Best$ to be $newBest$;
    end if
    for all fit pairs (p_1, p_2) in $population$ **do**
      Crossover (p_1, p_2) to get children C_1 and C_2;
      Mutate C_1 and C_2;
      Add C_1 and C_2 into $population$;
      Remove (p_1, p_2) from $population$;
    end for
    $generation \leftarrow generation + 1$;
  until $generation > someThreshold$
  $Z \leftarrow Z + 1$;
until $Best$ is not improved
return $Best$

the learned model). In particular, a single parameter μ is used in our approach, whereas in existing algorithms [23,9], a parameter μ is used to select the value of ϵ (based on a false assumption of the BIC being concave), which in turn controls the degree of generalization. From a user point of view, it is hard to see the effect of having a different ϵ value since it controls whether two nodes are merged in the intermediate steps of the learning process.

Next, we discuss how candidate models with better fitness score are selected. Selection directs evolution towards better models by keeping good chromosomes and weeding out bad ones based on their fitness. Two standard selection strategies are applied. One is *roulette wheel selection.* Suppose f is the average fitness of a population. For each individual M in the population, we select f_M/f copies of M. The other is *tournament selection.* Two individuals are chosen randomly from the population and a tournament is staged to determine which one gets selected. The tournament is done by generating a random number r between zero and comparing it to a pre-defined number p (which is larger than 0.5). If r is smaller than p, the individual with a higher fitness score is kept. We refer the readers to [18] for discussion on the effectiveness of these selection strategies.

After selection, genetic operators like mutation and crossover are applied to the selected candidates. Mutation works by mapping a random node to a new number from the same group, i.e., merging the node with other nodes with the same last observation. For crossover, chromosomes in the current generation are randomly paired and two

children are generated to replace them. Following standard approaches [18], we adopt three crossover strategies.

- *One-point Crossover.* A crossover point is randomly chosen, one child gets its prefix from the father and suffix from the mother. Reversely for the other child.
- *Two-point Crossover.* Two crossover points are randomly chosen, which results in two crossover segments in the parent chromosomes. The parents exchange their crossover segments to generate two children.
- *Uniform Crossover.* One child gets its odd bit from father and even bit from mother. Reversely for the other child.

We remark that during mutation or crossover, we guarantee that only chromosomes representing valid DTMC models are generated, i.e., only two nodes with the same last observations are mapped to the same number (i.e., a state in the learned model).

The details of our GA-based algorithm is shown as Algorithm 2. Variable Z is the number of states in the learned model. We remark that the number of states in the learned model M is unknown in advance. However, it is at least the number of letters in alphabet Σ, i.e., when all nodes in $tree(\Pi)$ sharing the same last observation are merged. Since a smaller model is often preferred, the initial population is generated such that each of the candidate models is of size $|\Sigma|$. The size of the model is incremented by 1 after each round of evolution. Variable $Best$ records the fittest chromosome generated so far, which is initially set to be $null$ (i.e., the least fit one). At line 3, an initial population of chromosome with Z states are generated as discussed above. The loop from line 5 to 17 then lets the population evolve through a number of generations, during which crossover and mutations take place. At line 18, we then increase the number of states in the model in order to see whether we can generate a fitter chromosome. We stop the loop from line 2 to 19 when the best chromosome is not improved after increasing the number of states. Lastly, the fittest chromosome $Best$ is decoded to a DTMC and presented as the learned model.

Example We use an example to illustrate how the above approach works. For simplicity, assume we have the following collection of executions $\Pi = \{\langle aacd\rangle, \langle abd\rangle, \langle acd\rangle\}$ from the model shown in Figure 1. There are in total 10 prefixes of these execution (including the empty string). As a result, the tree $tree(\Pi)$ contains 10 nodes. Since the alphabet $\{a, b, c, d\}$ has size 4, the nodes (except the root) are partitioned into 4 groups so that all nodes in the same group have the same last observation. The initial population contains a single model with 4 states, where all nodes in the same groups are mapped into the same state. After one round of evolution, models with 5 states are generated (by essentially splitting the nodes in one group to two states) and evaluated with the fitness function. The evolution continues until the fittest score does not improve anymore when we add more states.

4.3 Learn from Single Execution

In the following, we describe our GA-based learning if there is only one system execution. Recall that we are given a single long system observation α in this setting. The goal is to identify the *shortest dependent history memory* that yields the most precise

probability distribution of the system's next observation. That is, we aim to construct a part of $tree(\alpha)$ which transforms to a "good" DTMC. A model thus can be defined as an assignment of each node in $tree(\alpha)$ to either true or false. Intuitively, a node is assigned true if and only if it is selected to predict the next observation, i.e., the corresponding suffix is kept in the tree which later is used to construct the DTMC model. A chromosome (which encodes a model) is thus in the form of a sequence of boolean variable $\langle B_1, B_2, \cdots, B_m \rangle$ where B_i represents whether the i-th node is to be kept or not. We remark that not every valuation of the boolean variables is considered a valid chromosome. By definition, if a suffix π is selected to predict the next observation, all suffixes of π are not selected (since using a longer memory as in π predicts better) and therefore their corresponding value must be false. During mutation and crossover, we only generate those chromosomes satisfying this condition so that only valid chromosomes are generated.

A chromosome defined above encodes a part of $tree(\alpha)$, which can be transformed into a DTMC following the approach in [31]. Let M be the corresponding DTMC. The fitness function is defined similarly as in Section 4.2. We define the fitness function of a chromosome as $log(Pr_M(\alpha)) - \mu \times |M| \times log(|\alpha|)$ where $Pr_M(\alpha)$ is the probability of exhibiting α in M, μ is a constant that controls the weight of model size, and $|\alpha|$ is the size of the input execution. Mutation is done by randomly selecting one boolean variable from the chromosome and flip its value. Notice that afterwards, we might have to flip the values of other boolean values so that the chromosome is valid. We skip the discussion on selection and crossover as they are the same as described in Section 4.2.

We remark that, compared to existing algorithms in learning models [23,9,24], it is straightforward to argue that the GA-based approaches for model learning do not rely on the assumption needed for BIC. Furthermore, the learned model improves monotonically through generations.

5 Empirical Study

The above mentioned learning algorithms are implemented in a self-contained tool called ZIQIAN (available at [37], approximately 6K lines of Java code). In this work, since the primary goal of learning the models is to verify properties over the systems, we evaluate the learning algorithms by checking whether we can reliably verify properties based on the learned model, by comparing verification results based on the learned models and those based on the actual models (if available). All results are obtained using PRISM [22] on a 2.6 GHz Intel Core i7 PC running OSX with 8 GB memory. The constant μ in the fitness function of learning by GA is set to 0.5.

Our test objects can be categorized in two groups. The first group contains all systems (*brp*, *lse*, *egl*, *crowds*, *nand*, and *rsp*) from the PRISM benchmark suite for DTMCs [21] and a set of randomly generated DTMC models (*rmc*) using an approach similar to the approach in [36]. We refer the readers to [21] for details on the PRISM models as well as the properties to be verified. For these models, we collect multiple executions. The second group contains two real-world systems, from which we collect a single long execution. One is the probabilistic boolean networks (*PBN*), which is a modeling framework widely used to model gene regulatory networks (GRNs) [34]. In

PBN, a gene is modeled with a binary valued node and the interactions between genes are expressed by Boolean functions. For the evaluation, we generate random PBNs with 5, 8 and 10 nodes respectively using the tool ASSA-PBN [25]. The other is a real-world raw water purification system called the Secure Water Testbed (*SWaT*) [35]. *SWaT* is a complicated system which involves a series of water treatments like ultrafiltration, chemical dosing, dechlorination through an ultraviolet system, etc. We regard *SWaT* as a representative complex system for which learning is the only way to construct a model. Our evaluation consists of the following parts (all models as well as the detailed results are available at [38]).

We first show that assumptions required by existing learning algorithms may not hold, which motivates our proposal of GA-based algorithms. Existing learning algorithms [23,9] require that the BIC score is a concave function of ϵ in order to select the best ϵ value which controls the degree of generalization. Figure 2 shows how the absolute value of BIC scores ($|BIC|$) of representative models change with ϵ. It can be observed that this assumption is not satisfied and ϵ is not controlling the degree of generalization nicely. For example, the $|BIC|$ (e.g., for *brp*, *PBN* and *egl*) fluctuate with ϵ. Besides, we observe climbings of $|BIC|$ for *lse* when ϵ increases, but droppings for *crowds*, *nand* and *rsp*. What's worse, in the case (e.g., *PBN*) of learning from a single execution, if the range of ϵ is selected improperly, it is very likely that an empty model (a tree only with root $\langle\rangle$) is learned.

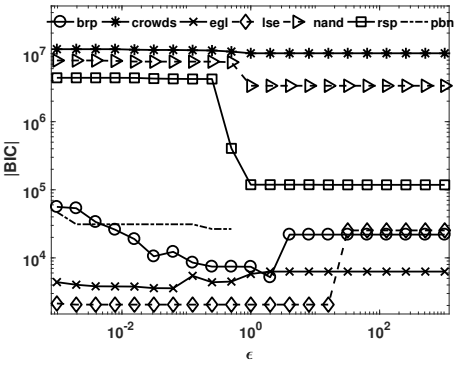

Fig. 2: How the absolute values of BIC score change over ϵ.

Second, how fast does learning converge? In the rest of the section, we adopt absolute relative difference (ARD) as a measure of accuracy of different approaches. The ARD is defined as $|P_{est} - P_{act}|/P_{act}$ between the precise result P_{act} and the estimated results P_{est}, which can be obtained by AA, GA as well as SMC. A smaller ARD implies a better estimation of the true probability. Figure 3 shows how the ARD of different systems change when we gradually increase the time cost from 30 seconds to 30 minutes by increasing the size of training data. We remark that some systems (*brp*, *egl*, *lse*) are not applicable due to different reasons. We can observe that GA converges faster and better than AA. In general, both AA and GA converges to relatively accurate results when we are given sufficient time. But there are also cases of fluctuation of ARD, which is problematic, as in such cases, we would not know which result to trust (given the different verification results obtained with different number of sampled executions), and it is hard to decide whether we have gathered enough system executions for reliable verification results.

Third, how accurate can learning achieve? We compare the accuracy of AA, GA, and SMC for benchmark systems given the same amount of time in Fig. 4. We remark that due to the discrimination of system complexity (state space, variable number/type,

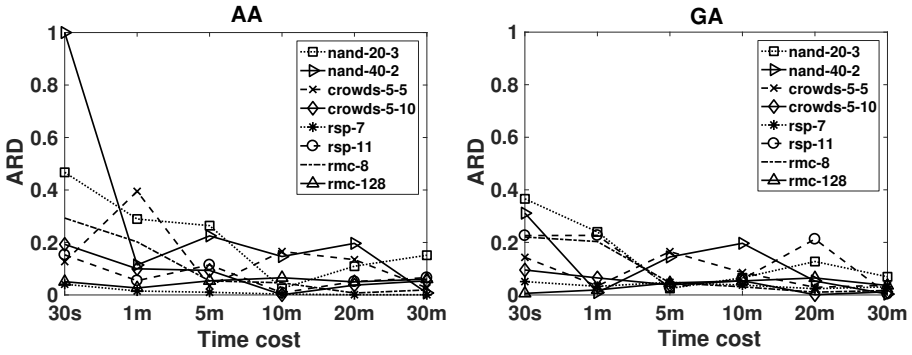

Fig. 3: Convergence of AA and GA over time. The numbers after the system of legends are one kind of system configuration.

etc.), different systems can converge in different speed. For SMC, we adopt the statistical model checking engine of PRISM and select the confidence interval method. We fix confidence to 0.001 and adjust the number of samples to adjust time cost. We have the following observations based on Fig. 4. Firstly, for most systems, GA results in more accurate results than AA given same amount of time. This is especially true if sufficient time (20m or 30m) are given. *However, it should be noticed that SMC produces significantly more accurate results*. Secondly, we observe that model learning works well if the actual model contains a small number of states. Cases like random models with 8 states (rmc-8) are good examples. For systems with more states, the verification results could deviate significantly (like nand-20-3, rsp-11).

Among our test subjects, *PBN* and *SWaT* are representative systems for which manual modelling is extremely challenging. Furthermore, SMC is not applicable as it is infeasible to sample the executions many times for these systems. We evaluate whether we can learn precise models in such a scenario. Note that since we do not have the actual model, we must define the preciseness of the learned model without referring to the actual model. For *PBN*, following [34], we use mean squared error (MSE) to measure how precise the learned models are. MSE is computed as follows: $MSE = \frac{1}{n}\sum_{i=1}^{n}(\hat{Y}_i - Y_i)^2$ where n is the number of states in *PBN* and Y_i is the steady-state probabilities of the original model and $\hat{Y}_i$ is the corresponding steady-state probabilities of the learned model. We remark that the smaller its value is, the more precise the learned model is. Table 1 shows the MSE of the learned models with for *PBN* with 5, 8, and 10 nodes respectively. Note that AA and GA learn the same models and thus have the same MSE, while GA always consumes less time. We can observe the MSEs are very small, which means the learned models of *PBN* are reasonably precise.

For the *SWaT* system, we evaluate the accuracy of the learned models by comparing the predicted observations against a set of test data collected from the actual system. In particular, we apply steady-state learning proposed in [9] (hereafter SL) and GA to learn from executions of different length and observe the trends over time. We select 3 critical sensors in the system (out of 50), named $ait502$, $ait504$ and $pit501$, and learn models on how the sensor readings evolve over time. During the experiments, we find it very difficult to identify an appropriate ϵ for SL in order to learn a non-empty useable

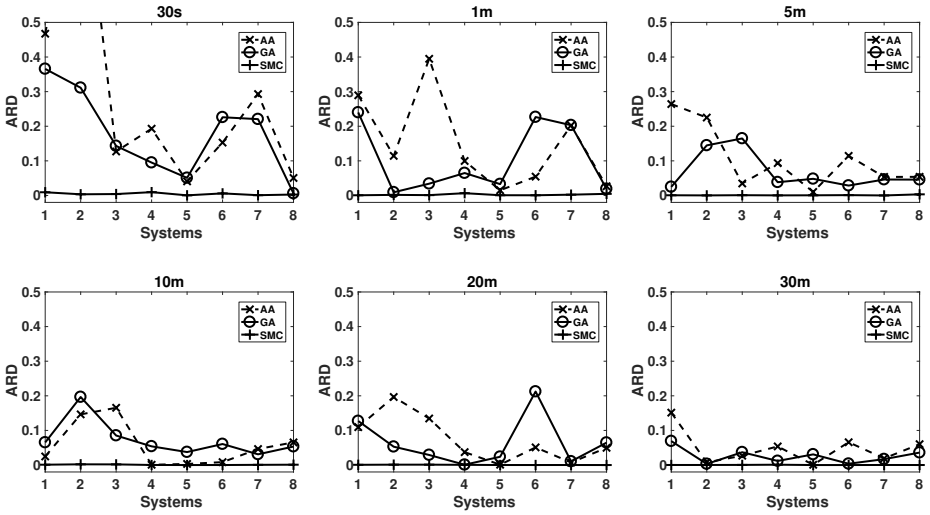

Fig. 4: The comparison of accuracy of AA, GA, and SMC given same amount of time, which varies from 30 seconds to 30 minutes. The horizontal-axis point represents a benchmark system with certain configuration in Fig. 3.

model. Our GA-based approach however does not have such problem. Eventually we managed to identify an optimal ϵ value and both SL and GA learn the same models given the same training data. A closer look at the learned models reveals that they are all first-order Markov chains. This makes sense in the way that sensor readings in the real *SWaT* system vary slowly and smoothly. Applying the learned models to predict the probability of the test data (from another day with length 7000), we observe a very good prediction accuracy. We use the average prediction accuracy for each observation $\bar{P}_{obs} = P_{td}^{1/|td|}$, where td is the test data and $|td|$ is its length, to evaluate how good the models are. In our experiment, the average accuracy of prediction for $ait502$ and $pit501$ is over 0.97, and the number is 0.99 for $ait504$, which are reasonably precise.

Last, there are some potential problems that may render learning ineffective. One of them is the known problem of rare-events. For *brp* system, the probability of satisfying the given properties are very small. As a result, a system execution satisfying the property is unlikely to be observed and learned from. Consequently, the verification results based on the learned models are 0. It is known that SMC is also ineffective for these properties since it is also based on random sampling. Besides, learning doesn't work when the state space of underlying system is too large or even infinite. If there are too many variables to observe (or when float/double typed variables exist), which induces a very large state space, learning will become infeasible. For example, to verify the fairness property of *egl* protocol, we need to observe dozens of integer variables. Our experiment suggests that AA and GA take unreasonable long time to learn a model, e.g., more than days. In order to apply learning in this scenario, we thus have to apply abstraction on the sampled system executions and learn from the abstract traces. Only by

Table 1: Results of *PBN* steady-state learning.

# nodes	# states	trajectory size ($\times 10^3$)	time cost(s) SL	time cost(s) GA	MSE ($\times 10^{-7}$)	# nodes	# states	trajectory size ($\times 10^3$)	time cost(s) SL	time cost(s) GA	MSE ($\times 10^{-7}$)
5	32	5	37.28	6.37	36.53	8	256	5	29.76	2.36	1.07
		15	161.57	53.49	15.21			15	105.87	26.4	0.03
		25	285.52	182.97	6.04			25	197.54	73.92	0.37
		35	426.26	348.5	7.75			35	310.87	122.61	0.94
		45	591.83	605.1	5.74			45	438.09	429.81	0.78
		50	673.55	767.7	4.28			50	509.59	285.66	0.34
10	1024	5	902.69	266.74	1.78	10	1024	15	5340.54	2132.68	0.61
		10	2772.56	1010.16	1.01			20	8477.24	3544.82	0.47

doing so, we are able to reduce the learning time significantly (in seconds) and successfully verified the *egl* protocol by learning. However, how to identify the right level of abstraction is highly non-trivial in general and is to be investigated in the future. What's more, there are other complications which might make model learning ineffective. For the *lse* protocol, the verification results based on the learned models may deviate from actual result for properties that show the probability of electing a leader in L rounds, with a different value for L. While the actual result 'jumps' as L increases, the result based on the learned model is smooth and deviates from actual results significantly when L is 3, 4 or 5, while results based on SMC are consistent with the actual results.

6 Conclusion and Related Work

In this work, we investigate the validity of model learning for the purpose of PMC. We propose a novel GA-based approach to overcome limitations of existing model learning algorithms and conducted an empirical study to systematically evaluate the effectiveness and efficiency of all these model learning approaches compared to statistical model checking over a variety of systems. We report their respective advantages and disadvantages, potential applications and future direction to improve.

This work is inspired by the work on comparing the effectiveness of PMC and SMC [40] and the line of work on adopting machine learning to learn a variety of system models (e.g., DTMC, stationary models and MDPs) for system model checking, in order to avoid manual model construction [23,9,24]. Existing learning algorithms are often based on algorithms designed for learning (probabilistic) automata, as evidenced in [31,30,7,17,8,1]. Besides the work in [23,9,24] which have been explained in detail, this work is also related to the work in [32], which learns continuous time Markov chains. In addition, in [6], learning algorithms are applied in order to verify Markov decision processes, without constructing explicit models. Our proposal on adopting genetic algorithms is related to work on applications of evolutionary algorithms for system analysis. In [14], evolutionary algorithm is integrated to abstraction refinement for model checking. This work is remotely related to work on SMC [41,33], some recent work on extending SMC to unbounded properties [39,29]. Lastly, our work uses the PRSIM model checker as the verification engine [22] and the case studies are taken from various practical systems and protocols including [16,15,28,27,26,19,25].

References

1. Angluin, D.: Learning regular sets from queries and counterexamples. Information and computation 75(2), 87–106 (1987)
2. Angluin, D.: Identifying languages from stochastic examples (1988)
3. Baier, C., Katoen, J.P., et al.: Principles of model checking, vol. 26202649. MIT press Cambridge (2008)
4. Bauer, A., Leucker, M., Schallhart, C.: Monitoring of real-time properties. In: FSTTCS 2006: Foundations of Software Technology and Theoretical Computer Science, pp. 260–272. Springer (2006)
5. Bianco, A., De Alfaro, L.: Model checking of probabilistic and nondeterministic systems. In: Foundations of Software Technology and Theoretical Computer Science. pp. 499–513. Springer (1995)
6. Brázdil, T., Chatterjee, K., Chmelík, M., Forejt, V., Křetínskỳ, J., Kwiatkowska, M., Parker, D., Ujma, M.: Verification of markov decision processes using learning algorithms. In: Automated Technology for Verification and Analysis, pp. 98–114. Springer (2014)
7. Carrasco, R.C., Oncina, J.: Learning stochastic regular grammars by means of a state merging method. In: Grammatical Inference and Applications, pp. 139–152. Springer (1994)
8. Carrasco, R.C., Oncina, J.: Learning deterministic regular grammars from stochastic samples in polynomial time. Informatique théorique et applications 33(1), 1–19 (1999)
9. Chen, Y., Mao, H., Jaeger, M., Nielsen, T.D., Larsen, K.G., Nielsen, B.: Learning markov models for stationary system behaviors. In: NASA Formal Methods, pp. 216–230. Springer (2012)
10. Clarke, E.M., Grumberg, O., Peled, D.: Model checking. MIT press (1999)
11. Clarke, E.M., Zuliani, P.: Statistical model checking for cyber-physical systems. In: Automated Technology for Verification and Analysis, pp. 1–12. Springer (2011)
12. Dyer, D.W.: Watchmaker framework for evolutionary computation. `http://watchmaker.uncommons.org`
13. Havelund, K., Roşu, G.: Synthesizing monitors for safety properties. In: Tools and Algorithms for the Construction and Analysis of Systems, pp. 342–356. Springer (2002)
14. He, F., Song, X., Hung, W.N., Gu, M., Sun, J.: Integrating evolutionary computation with abstraction refinement for model checking. Computers, IEEE Transactions on 59(1), 116–126 (2010)
15. Helmink, L., Sellink, M.P.A., Vaandrager, F.W.: Proof-checking a data link protocol. Springer (1994)
16. Herman, T.: Probabilistic self-stabilization. Information Processing Letters 35(2), 63–67 (1990)
17. De la Higuera, C.: Grammatical inference, vol. 96. Cambridge University Press Cambridge (2010)
18. Holland, J.H.: Adaptation in Natural and Artificial Systems. MIT Press, Cambridge, MA, USA (1992)
19. Itai, A., Rodeh, M.: Symmetry breaking in distributed networks. Information and Computation 88(1), 60–87 (1990)
20. Kermorvant, C., Dupont, P.: Stochastic grammatical inference with multinomial tests. In: Grammatical Inference: Algoritihms and Applications, pp. 149–160. Springer (2002)
21. Kwiatkowska, M., Norman, G., Parker, D.: The PRISM benchmark suite. In: Proc. 9th International Conference on Quantitative Evaluation of SysTems (QEST'12). pp. 203–204. IEEE CS Press (2012)
22. Kwiatkowska, M., Norman, G., Parker, D.: Prism: Probabilistic symbolic model checker. In: Computer performance evaluation: modelling techniques and tools, pp. 200–204. Springer (2002)

23. Mao, H., Chen, Y., Jaeger, M., Nielsen, T.D., Larsen, K.G., Nielsen, B.: Learning probabilistic automata for model checking. In: Quantitative Evaluation of Systems (QEST), 2011 Eighth International Conference on. pp. 111–120. IEEE (2011)
24. Mao, H., Chen, Y., Jaeger, M., Nielsen, T.D., Larsen, K.G., Nielsen, B.: Learning markov decision processes for model checking. arXiv preprint arXiv:1212.3873 (2012)
25. Mizera, A., Pang, J., Yuan, Q.: ASSA-PBN: a tool for approximate steady-state analysis of large probabilistic Boolean networks. In: Proc. 13th International Symposium on Automated Technology for Verification and Analysis. LNCS, Springer (2015), available at http://satoss.uni.lu/software/ASSA-PBN/
26. Norman, G., Parker, D., Kwiatkowska, M., Shukla, S.: Evaluating the reliability of nand multiplexing with prism. Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on 24(10), 1629–1637 (2005)
27. Norman, G., Shmatikov, V.: Analysis of probabilistic contract signing. Journal of Computer Security 14(6), 561–589 (2006)
28. Reiter, M.K., Rubin, A.D.: Crowds: Anonymity for web transactions. ACM Transactions on Information and System Security (TISSEC) 1(1), 66–92 (1998)
29. Rohr, C.: Simulative model checking of steady state and time-unbounded temporal operators. In: Transactions on Petri Nets and Other Models of Concurrency VIII, pp. 142–158. Springer (2013)
30. Ron, D., Singer, Y., Tishby, N.: On the learnability and usage of acyclic probabilistic finite automata. In: Proceedings of the eighth annual conference on Computational learning theory. pp. 31–40. ACM (1995)
31. Ron, D., Singer, Y., Tishby, N.: The power of amnesia: Learning probabilistic automata with variable memory length. Machine learning 25(2-3), 117–149 (1996)
32. Sen, K., Viswanathan, M., Agha, G.: Learning continuous time markov chains from sample executions. In: Quantitative Evaluation of Systems, 2004. QEST 2004. Proceedings. First International Conference on the. pp. 146–155. IEEE (2004)
33. Sen, K., Viswanathan, M., Agha, G.: Statistical model checking of black-box probabilistic systems. In: Computer Aided Verification. pp. 202–215. Springer (2004)
34. Shmulevich, I., Dougherty, E., Zhang, W.: From boolean to probabilistic boolean networks as models of genetic regulatory networks. Proceedings of the IEEE 90(11), 1778–1792 (2002)
35. SUTD: Secure water treatment testbed. http://itrust.sutd.edu.sg/research/testbeds/secure-water-treatment-swat/
36. Tabakov, D., Vardi, M.Y.: Experimental evaluation of classical automata constructions. In: Logic for Programming, Artificial Intelligence, and Reasoning, 12th International Conference, LPAR 2005, Montego Bay, Jamaica, December 2-6, 2005, Proceedings. pp. 396–411 (2005), http://dx.doi.org/10.1007/11591191_28
37. Wang, J.: ZIQIAN. https://bitbucket.org/jingyi_wang/ziqian_develop
38. Wang, J.: ZIQIAN evaluation. https://bitbucket.org/jingyi_wang/ziqian_evaluation
39. Younes, H.L., Clarke, E.M., Zuliani, P.: Statistical verification of probabilistic properties with unbounded until. In: Formal Methods: Foundations and Applications, pp. 144–160. Springer (2011)
40. Younes, H.L., Kwiatkowska, M., Norman, G., Parker, D.: Numerical vs. statistical probabilistic model checking. International Journal on Software Tools for Technology Transfer 8(3), 216–228 (2006)
41. Younes, H.L., Simmons, R.G.: Probabilistic verification of discrete event systems using acceptance sampling. In: Computer Aided Verification. pp. 223–235. Springer (2002)